

# Let's Talk AI with Moshe Y. Vardi

Moshe Y. Vardi<sup>1</sup> and Barbara Steffen<sup>2</sup>

<sup>1</sup> Rice University, Department of Computer Science, USA,  
vardi@rice.edu

<sup>2</sup> METAFrame Technologies GmbH,  
barbara.steffen@metaframe.de

*"My left brain is doing research in AI. My right brain is in a panic: What is going to happen? What will be the impact on society?"*

---

## The Interviewee - Moshe Y. Vardi



**My Personal AI Mission:**  
AI for the public good.

---

## My Takes on AI

**Artificial Intelligence:** The effort to simulate human intelligence by machines.

**Trust:** Reliance in risk-accepting situations.

**Explainability:** I do not need explanations. I need justifications.

**Essential Elements of Human Capabilities:** Humans are intrinsically sacred.

## The Interview

**Barbara** *Today I have the pleasure of interviewing Professor Moshe Vardi. Please introduce yourself and your relationship to AI.*

**Moshe Y.** My name is Moshe Vardi. I'm a professor of computer science at Rice University, where I've been for about 30 years. I find myself in a bit of a dilemma these days. My left brain is engaged in research in computer science, algorithms, and computational logic [1]. My right brain, on the other hand, constantly questions the left brain: "What are you doing?" [8]. This is a fitting description of my current state. So how does this relate to artificial intelligence? Well, I've been intensely interested in AI since 2011. A significant milestone in AI at that time was IBM Watson, a question-answering system. Watson competed in a TV game called Jeopardy, where players had to be knowledgeable about popular culture and history. Watson won decisively against two champion players, which was a big deal [10]. If you look at AI milestones, there was another significant one in the late '90s when IBM's Deep Blue beat Kasparov at chess [9]. This was a remarkable achievement that had been expected for many years. Watson was a surprise. After Watson's victory, I began to think that the realm of AI might be closer than we think. I started pondering the potential consequences, and the more I thought about it, the more concerned I became. So, over the past decade, I've been somewhat of a Cassandra, warning about the potential negative effects of technology on society [5]

**Barbara** *In your research on artificial intelligence, what specific challenges are you currently tackling?*

**Moshe Y.** In my technical research [4], I'm trying to bridge what people often refer to as System 1 and System 2 [3]. In a sense, our head houses two different brains. The back part, sometimes called the lizard brain, is something we share with almost all animals. It is responsible for recognizing danger, seeking food, and other basic survival instincts. But in the forebrain, we have developed the capacity for language and abstract reasoning, often referred to as slow thinking, as opposed to the fast thinking of the lizard brain. This is System 1 and System 2, in the terminology of the recently departed Daniel Kahneman. The big question for me is how to integrate these two systems in AI. Much of the progress in machine learning over the last decade has been in System 1. But we humans seamlessly transition between the two systems. We do not consciously think about switching from System 1 to System 2. For example, while driving, we recognize a stop sign (System 1), and then we decide to stop because that's the rule (System 2). The challenge for AI today is to integrate these two systems. This is partly where my research is focused today. Some people refer to it as neural-symbolic reasoning. Neural refers to neural nets that are the leading technology for machine learning, and symbolic refers to how we think of language and logic.

**Barbara** *Would you like AI to operate like humans? Or would you prefer AI to operate more at the System 2 level? I keep hearing that humans too often operate at the System 1 level, especially in situations where they would be better off operating at the System 2 level. Now, with AI, we may have the opportunity to nudge towards the preferred system for different contexts.*

**Moshe Y.** We all know that there are instances when we should have relied more on our rational thinking than our instinctive thinking. Nevertheless, when it comes to technology, we often draw inspiration from nature. Consider flying, as an example. Birds have wings and they flap them to fly. Early human attempts at flight tried to mimic this, but we eventually discovered that fixed

"I don't use [ChatGPT] because, for me, writing is thinking. I don't want to give that up."

wings and an engine were a more effective method. We drew inspiration from birds, but we didn't replicate their exact method. Similarly, while we can draw inspiration from the brain, we don't have

to mimic it exactly in AI. I hope that if we build intelligence, it won't be driven by instincts as much as we humans sometimes are. Evolution, often referred to as the blind watchmaker, doesn't always create what we might call a rational design.

**Barbara** *And when it comes to the use of AI, what role do you think trust plays, and what measures do you think we should take to ensure the ethical use of AI in the future?*

**Moshe Y.** The issue of trust in technology applies to almost all technology. We entrust our lives to technology every day. When you drive a car, you trust that it won't suddenly malfunction. When you board a plane, you trust that it's safe. A few weeks ago, I received my COVID booster shot, trusting that it was safe. We live in a world where we rely on many things that we don't fully understand. I understand the basic principle of an airplane, but I'm not an aeronautical engineer. I don't fully understand how it works, but I trust the system. Building that trust requires many elements, one of which is a responsible industry. One of the main challenges for our society is that the major players in our economy are corporations, and their primary goal is to make profits. They need, however, to understand that they also have a responsibility to society. For example, Boeing lost a lot of trust when it became clear that they had cut corners with the 737 MAX, leading to the deaths of hundreds of people. This caused significant reputational damage to Boeing, and only time will tell what impact it will have on their future sales. Companies need to be socially responsible, not just to avoid bad publicity and loss of profits, but also to uphold ethical standards [6]. This is particularly challenging in the tech world. Many tech companies claim to be responsible, but their actions don't always align with their words. For instance, Google has faced criticism for dismissing employees who raised ethical concerns about their technology. This has led to a loss of trust in the company, which was founded under the motto "Don't be evil." Facebook has faced similar issues. The problem is that most of us don't have the expertise to evaluate the

safety of these technologies ourselves, so we must trust the producers. But if the industry continues to lose the trust of society, people will start to lose trust in the technology itself.

**Barbara** *What do you think about the challenge that even if I don't trust the system, e.g. ChatGPT, if I don't adopt it into my daily work routine, I will be less productive than others who use it and start to fall behind. So, do we really have a choice?*

**Moshe Y.** Not really. Take smartphones, for example. I resisted getting a smartphone for a long time, but it's almost impossible to function in our society without one. I first resisted having a cell phone, then I held onto my flip phone as long as I could. But eventually, I had no choice but to get a smartphone. Now, I depend on it. There are places, like in China, where you can't survive without a smartphone. Even here, many services require a smartphone. For example, to confirm a doctor's appointment, you often must click on a link sent to your phone. Without a smartphone, this becomes impossible. So, it's not so much a personal choice anymore as it is a societal adaptation to technology. ChatGPT isn't quite at the level of necessity like a smartphone yet. I don't use it because, for me, writing is thinking. I don't want to give that up. It's like if everywhere you went, someone offered you a wheelchair to get around. You might think, "Do I really want that? If I use it all the time, I'll eventually lose my ability to walk. Is that a good thing?" You might decide to keep walking, even if it's less convenient. It's not always about being more productive. But we may reach a point where AI technology becomes such a basic part of our lives that we can't resist it. Often, industries try to frame these issues as matters of personal choice. For example, with sustainability, we're encouraged to recycle. I'm an avid recycler, but I'm skeptical that plastic recycling does any good. By framing these as matters of personal choice, industries can avoid regulation and systemic change. They tell us, "You choose. You decide." But can you really decide not to buy plastic products today? It's almost impossible to avoid plastic. So, by trying to frame systemic societal issues as matters of personal choice, industries avoid taking responsibility.

**Barbara** *Think about the future capabilities of AI on a scale of 1 to 10, with 1 standing for artificial intelligence systems we know today, such as ChatGPT, and 10 standing for so-called artificial general intelligence systems that surpass human capabilities. Where do you think we are headed?*

**Moshe Y.** Predicting the direction of technology is extremely difficult. As scientists, we may be able to see a little further, but even then, we can't predict

---

"I'm a materialist. I believe the brain is a machine. It's a biological machine with many neurons, and evolution has had billions of years to refine its design. We have proof that human-level intelligence can exist in a relatively small box. I wouldn't bet against our ability to achieve this with non-biological technology."

everything. Take ChatGPT, for example. It has given us the ability to generate human-level text. This happened about a year ago. If you had asked me two years ago when this would happen, I wouldn't have guessed it would be so soon. Very few people predicted it. We simply scaled up some models, and suddenly, at a certain point, they started generating human-like text. We don't fully understand how this happened. We don't understand AGI or human intelligence enough to predict when we'll achieve it. Some people suggest that once we reach 10 trillion parameters in a model, it will be as intelligent as a human. I'm skeptical of such predictions. Nevertheless, I'm a materialist. I believe the brain is a machine. It's a biological machine with many neurons, and evolution has had billions of years to refine its design. We have proof that human-level intelligence can exist in a relatively small box. I wouldn't bet against our ability to achieve this with non-biological technology. We really don't know, however, what will happen when we have something so intelligent. Will it just sit there and wait for instructions, or will it develop its own motivations? We don't know, and that's why I believe we need to slow down and consider the consequences before we continue to develop more advanced AI. When we developed nuclear weapons and nuclear energy, we immediately recognized the potential dangers and put regulations in place. I believe we need a similar approach for AI. The argument from the industry is that regulation will stifle innovation. But maybe we need to slow down risky innovation. Innovation is a means to an end, not a goal in itself. The goal should be the public good and human welfare [7].

**Barbara** *And when it comes to the future, what do you personally think, are we heading towards a utopian or dystopian future?*

**Moshe Y.** I consider myself a realist. I believe that technology can do amazing things. For example, we developed a vaccine for COVID in just one year.

So, we had a terrible pandemic. Many millions of people died. The last time we experienced something of this magnitude was 100 years ago, the Spanish flu. Now we realize that living through a pandemic is a horrific experience. The death toll is so high that it's difficult for our brains to comprehend. In the United States, we had close to a million and a half casualties. It's really hard to grasp the enormity of the death toll. But it could have been much worse. The two technologies that saved us are mRNA vaccines and the internet. We had a vaccine within a year. When the pandemic started, I told people it would be five years before there is a vaccine. The traditional method of creating a vaccine, which involves using a weakened strain of the virus, would have taken five years. And I think five years of a pandemic would have been disastrous. The internet also saved us. Suddenly, at the drop of a hat, I could work from home. Students could study from home. In the spring of 2020, we switched overnight to teaching online. Many people started working from home, and I started teaching from home. You could even have doctor appointments from home. Not everything, but many things can be done from home.

Imagine if we didn't have these technologies. The economic damage would have been much worse. There was severe economic damage due to the pandemic, and

this economic damage translates to loss of life. Economists have calculated this. But the economic damage would have been much, much worse. Some people would say, I have no choice, I must go back to work. And the death toll would have been much higher. So, the internet saved us. Social media helped us stay in touch. We are facing social isolation, which is also a known health risk. At my university, there was a suicide by a graduate student, which we believe happened because of social isolation. We had to make significant efforts to combat social isolation, especially among students.

But it's the same internet and social media that also created polarization in society. So, the story of technology has always been a double-edged sword. This has always been the case. Think about it, what is the most fundamental technology? Fire. If you look at pictures of early humans, they had large jaws. They needed to chew raw food. Once we discovered fire, we had cooked food. We no longer needed large jaws. In fact, our jaws shrunk faster than our teeth. We still have too many teeth. Our jaws are too small. We have a lot of dental problems because we have the same number of teeth but smaller jaws. Today, fire is incredibly important to us. Most cars still use fire to run. And power plants use fire to create power. But people also die from fire. So, the human quest, the human journey, has always been the story of humanity dealing with technology. We benefit from it, and it comes with risk. And we try to walk on the edge of benefiting from the technology, while minimizing and mitigating risks. I think the same will be true with AI. But we just need to learn from the experience. For example, consider fire. Fire is leading to climate change. Climate change is caused mostly by carbon, putting more carbon in the atmosphere. How do you get carbon in the atmosphere? You burn carbon, and CO<sub>2</sub> goes into the atmosphere. So, the climate is burning because we use fire. But fire is the technology that brought us where we are; we are today immensely richer than people over 250 years ago, before the Industrial Revolution. It is all based on fire. We have learned to harness carbon-based energy. And now the planet is in danger. So, we need to do a better job of anticipating the downside of technology. But it is not clear that scientists and engineers are the best people to anticipate the long-term consequences. I would take the writers of Black Mirror. Here, is a funny anecdote. There was a review, I think last year, in in "Wired" magazine (a very techie magazine) of the most recent season of Black Mirror. And the reviewer complained. He wrote "Black Mirror used to be dystopian. Now it's like a reality show."

So, we need to think about what can go wrong? And the answer is we should try to anticipate it. But, very often, by the time we realize it, the damage has already happened. So, now we try to take social media and try to regulate it. It's going to be very, very difficult. So, can we basically slow down innovation because we need to anticipate risk? Can we do it in a more deliberate way? The problem for society right now is that we are a capitalism-driven society. And it means that we just pursue profits. And we do it in a reckless way. Then we discover the damage later. Sometimes we can mitigate it later. And sometimes it

might be too late for that. And that's our problem right now. Look for example at what happened with climate change. If you have a coal economy based on carbon, a lot of parties would say, no, no, no. There's no climate change. It's caused because of sunspots or something else. Because now it's going to cost them a lot of money to change.

**Barbara** *Reflecting on the last few days, what was a particularly interesting insight from another discipline?*

**Moshe Y.** The most interesting insight was about explainable AI. One insight I got from a talk this morning on explainable AI is that we really don't care about explanation. What we really care about is justification. Is the recommendation justified? I don't care if I understand why it was given, but is it justified? If you recommend something to me, I'm going to ask you for an explanation. But I won't be satisfied with simply getting an explanation. I want to know; do I accept the explanation? Has Barbara given me justification? And so that's an insight that came to me this morning. But I think the most important insight is the importance of multidisciplinary conversations. So, the fact that computer scientists are interacting with lawyers and giving feedback to lawyers about the computer-science perspective and letting those different perspectives inform our own thinking. The same goes for conversations with philosophers and political scientists. These are very, very valuable conversations that do not normally happen. Because science in general has become so advanced that we each build our own little silo. We only talk to the people in our discipline because they understand our language. We don't have to go out of our way to explain something to people who don't understand our language. So, this is something that we really need to fight. And we need to have more meetings like AISoLA and bring different disciplines together. And for me that is the biggest insight. Really... from that point of view it's a remarkable experience.

"We simply scaled up some models, and suddenly, at a certain point, they started generating human-like text."

cal scientists. These are very, very valuable conversations that do not normally happen. Because science in general has become so advanced that we each build our own little silo. We only talk to the people in our discipline because they understand our language. We don't have to go out of our way to explain something to people who don't understand our language. So, this is something that we really need to fight. And we need to have more meetings like AISoLA and bring different disciplines together. And for me that is the biggest insight. Really... from that point of view it's a remarkable experience.

**Barbara** *That is great to hear! Is there a specific research question that needs to be urgently addressed? Or are you referring to a more general approach of looking at AI and its future development from an interdisciplinary perspective?*

**Moshe Y.** That tells me, for example, that if we're going to build a regulatory agency for AI, it has to be interdisciplinary. It should not just have AI experts in it. It should have lawyers. It should have philosophers. We need to bring together the different people who can shed light from different perspectives. Each one of us can shed light from a relatively narrow perspective. But if you really want to see the whole phenomenon, we need all these things. And we need interaction among these people. We need a conversation between them. The complex phenomenon we are dealing with here is the interaction of technology and society. But technologists don't understand enough about society. And people who do think about society, technology is a bit of a black box to them. So, we have

to look at that and say, OK, how do we bring in all the different perspectives? And how do we create trustworthiness, safety, and security? How do we better manage the interaction between technology and society?

**Barbara** *And from your personal perspective, what should be the AI vision for the coming years?*

**Moshe Y.** My vision for AI is to "slow down". Today, we have a very rich and powerful industry. And they're just rushing ahead. And I say just slow down. Let's talk about it. What technology are you pushing? Let's talk about the benefits. Let's talk about the risks. Let's talk about the consequences. We don't usually have that conversation about new technologies. Think about the standard process for new technologies. A company comes up with a new product. For example, Apple introduced the iPhone. Poof, we have an iPhone. And now we have smartphones, and we gave them to children. And history is not going to look kindly on us. We have a mental-health crisis among young people. It is hard to find any cause for that other than this technological combination of smartphones and social media. Now we have a rising suicide rate among children. And the only credible explanation is the combination of smartphones and social media. And who is responsible? We can blame Big Tech. But ultimately, we as a society have allowed this to happen. We allowed it to happen by letting innovation proceed without asking, "what will be the consequences"? To me, that is the critical question. How to stop technology from rushing ahead and presenting us with a fait accompli? How do we slow it down, think about consequences, and reflect on our responsibility? Who should be responsible? People talk about responsible AI. And I ask, what does it mean "responsible AI"? Who is responsible? Are you talking about the technology being responsible? Are the people being responsible? Should the industry be responsible? Who is responsible? And that's why I say, when people say responsible AI, I ask, what does it mean? Whom do you try to make more responsible?

---

"The problem for society right now is that we are a capitalism-driven society. And it means that we just pursue profits. And we do it in a reckless way. Then we discover the damage later."

**Barbara** *Isn't that the general problem, that we don't have a person or an institution that is or feels responsible for 2050, 2100, or even further into the future?*

**Moshe Y.** No, we don't. We believe that human history will move forward, but can we direct it? Communism was an attempt to say, well, we are going to steer society in the direction that we collectively think it should go. And it proved to be a disaster because the only way you can have that level of tight control is through incredibly coercive means. It is estimated that the death toll of the two Communist revolutions, the Soviet and the Chinese, was 100 million people. So, it's very risky to say that we should collectively decide where humanity should go. So, we have to be very modest about what we can do. But we should look at where it is going. What do we identify as risk? Let's think about risk in advance.



Let's have forward-looking laws. The legal system is way behind, I believe. We've let the technology run forward, and now we're trying to think, "OK, who should be liable for this?" We need to do some slowing down and serious thinking. I believe that technologists have proven to be irresponsible, just thinking about the technology and not thinking about the consequences and responsibility.

If we want responsible technology, then we need to put the responsibility on people. This is ethics. And the fundamental principle of ethics for technology should be that technologists need to think in advance about the consequences of the technology. And they need to ensure that this technology is ultimately for the benefit of humanity. This is where we have a huge deficit, I believe. We train people that business is about making money. That's how we train business students. (That's how we train students in the United States. I don't know how they get trained in other countries.) They teach them ethics, but ethics is usually taught in a very narrow sense. I've seen boards on ethics, where the basic message is, "don't cheat, don't lie, don't steal". And for engineers, well, "if you build a bridge, it should stand and not fall". That is kind of professional ethics in a nutshell: "Whatever you do, be honest, and your design should be good." But that does not address the societal consequences of technology [Ferreira and Vardi, 2021]. And I think that ethics comes from capitalism, where we just maximize profits. Put crassly, greed is good; it's OK for me to be greedy. But deeper ethics says "No"; it is not ok to be greedy. You must consider society. So, we must balance capitalism with ethics; what some people call "Moral Capitalism". We know that economic growth benefits society, but we need to find a way to do it in a societally responsible way. So far, we have not found a good way to do it. Humanity experimented with communism; it was a disaster. Everybody acknowledged that it was a disaster. Then we said, OK, forget it. It was a failure. So, let's have neoliberalism, which some people call "market fundamentalism." And now we're seeing the results of that, in the deeply polarized US and UK. So now we need to go back and ask, what is the right balance? And that's a fundamental challenge; in life, you're always searching for balance. You want to have some personal life. You want some progress. Progress as a student, as a professional, right? We talk about work-life balance. Is there a formula? There's no formula. We're trying to find a balance. When you raise children, you're faced with the issue of regulation. You're trying to regulate your children. How much should you regulate them? If you give them complete freedom, they will do crazy things. They might kill themselves. So, you say, no, no, I need to have some discipline. You do too much of that, you raise children who have not learned to make decisions, to take responsibility. So, we are all looking for balance, navigating between extremes. An ancient Asian wisdom is that everything has to be in balance. And my feeling is that one of the most fundamental questions for modern society is the balance between the market and the state. And we need to find a better balance than we have right now.

**Barbara** *Do you have anything to add?*

**Moshe Y.** For me, a big revelation that came out of my thinking about modern society over the last few years is the concept of the public good. And it turns out that if you go back to Jewish ethics from about 2,000 years ago, it is there in some sense. First, it says, if you are not for yourself, then who is going to be? First, you must take care of yourself. But if you're only for yourself, what kind of a person are you? And this is our challenge, right? I would say this is a conversation we need to have starting from a very early age, how to balance our own benefit versus our responsibility to society. And imagine you have someone like a politician, like Margaret Thatcher, who said, "there's no such thing as society". This is a horrendous, horrendous statement, OK? Because none of us can survive without society. It takes a village. We all are contributing and benefiting from society. Thatcher

"My vision for AI is to "slow down". Today, we have a very rich and powerful industry. And they're just rushing ahead. And I say just slow down. [...] Let's talk about it. What technology are you pushing? Let's talk about the benefits. Let's talk about the risks. Let's talk about the consequences."

basically said, just be greedy, don't care about society. It's a recipe for dystopia. And I've asked other people, I said, tell me about your education. When did you have a serious discussion about the public good and your responsibility for the public good? Most people say, no, this is not in my education. We never discussed the public good. And I think the public good is incredibly important. And especially right now, we have a planet that is burning. Without the public good, we're not going to survive. The public good is

the collective good of all of us. And that's something that our society has neglected. But also with regards to collective action towards the future we want as a society. At the end of the day, we're individuals. And each one of us has their own thinking about what needs to be done.

Consider the US Constitution, which to me is an amazing document. So, it starts with the first paragraph. It says the purpose of the United States is the public good. But how do we collectively agree on what the public good is? The answer is, it's very difficult. So, we try to get democracies and develop democratic systems for agreeing on what the public good is. We don't have one way to do democracy. The founding fathers of the United States, the US was the first republic that tried that, said, here is a proposal how to do democracy. This was the first constitution. It was the Constitution of the United States. Countries after that say, OK, yeah, that's not a bad idea. It's a set of rules that try to govern. How do we reach agreement? Because it's very difficult. Even in a family, two people, they must reach agreement. Just a couple, they must reach agreement. It's not always easy. Academic departments are kind of mini democracies. We must make decisions collectively. It is hard. Everybody pulls in a different direction. This is a kind of challenge to society, to find a way to contribute to the public good and find agreement on what is the public good. Nobody is against the public good. The devil is in the details. And we collectively must decide what is the public good and how do we collectively move towards the public good. It is very difficult.

**Barbara** *Especially since it is likely that everyone has a different idea of what is desired, what is to be expected, and what responsibility each individual should carry.*

**Moshe Y.** Indeed. But I think we need to start this conversation. I think we have not really started this conversation seriously. I'm looking at my education. I never participated in a serious conversation about it. We had classes in civics, and we learned about how the government works: the executive branch, the legislative branch, and the judicial branch.

It was very descriptive. But there was very little philosophy. What is the purpose of all these different branches and how they operate? What is the purpose? The purpose is to agree collectively on

what we think the public good is and agree on the actions we collectively need to take to promote the public good. This discussion was missing; it was not there. Now consider the people who graduated with this education. They may be able to recite what the different branches are, et cetera, but very little about the philosophy. And the philosophy is really the principle that should govern our life as a society. This is the difference between a civilized society and what people call the "Law of the Jungle". The Law of the Jungle is the mighty that takes what the mighty wants. In civilized society, we have a rule of law. What is the purpose of the law. You can go back to Hammurabi 4,000 years ago. What is the purpose of the law, asked Hammurabi? The purpose of the law is for the benefit of humanity, he answered, for the benefit of humanity.

"By framing [...] systemic societal issues as matters of personal choice, industries avoid taking responsibility."

**Barbara** *That's a perfect last sentence. Thank you, Moshe, for your time, your insights, and your perspective on the future.*

## References

1. Bansal, S., Li, Y., Tabajara, L. M., Vardi, M. Y., & Wells, A. (2023, October). Model Checking Strategies from Synthesis over Finite Traces. In International Symposium on Automated Technology for Verification and Analysis (pp. 227-247). Springer.
2. Ferreira, R., & Vardi, M. Y. (2021). Deep tech ethics: An approach to teaching social justice in computer science. In Proceedings of the 52nd ACM Technical Symposium on Computer Science Education, pp. 1041-1047.
3. Kahneman, D. (2011). Thinking, fast and slow. New York: Farrar, Straus and Giroux. ISBN: 9780374275631 0374275637
4. Prates, M., Avelar, P. H., Lemos, H., Lamb, L. C., & Vardi, M. Y. (2019, July). Learning to solve NP-complete problems: A graph neural network for decision TSP. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 33, No. 01, pp. 4731-4738).
5. Vardi, M. Y. (2015). Is information technology destroying the Middle Class? Communications of the ACM, 58(2), 5-5.
6. Vardi, M. Y. (2018). Are we having an ethical crisis in computing? Communications of the ACM, 62(1), 7-7.

7. Vardi, M. Y. (2023a). ACM for the Public Good. *Communications of the ACM*, 66(5), 5-5
8. Vardi, M. Y. (2023b). Computing, You Have Blood on Your Hands! *Communications of the ACM*, 67(1), 5-5.
9. [https://en.wikipedia.org/wiki/Deep\\_Blue\\_versus\\_Garry\\_Kasparov](https://en.wikipedia.org/wiki/Deep_Blue_versus_Garry_Kasparov)
10. [https://en.wikipedia.org/wiki/IBM\\_Watson](https://en.wikipedia.org/wiki/IBM_Watson)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

